

Present and Future Computing Requirements for the Dark Energy Survey

Scott Dodelson

Fermilab

NERSC BER Requirements for 2017
November 27, 2012
Rockville, MD

Dark Energy Survey

Fermilab, UIUC/NCSA, University of Chicago, LBNL , NOAO, University College London, University of Cambridge, University of Edinburgh, University of Portsmouth, University of Sussex, University of Nottingham, Instituto de Ciencias del Espacio, Institut de Fisica d'Altes Energies, Centro de Investigaciones Energeticas, Medioambientales y Tecnologicas, The University of Michigan, Observatorio Nacional, Centro Brasileiro de Pesquisas Fisicas, Universidade Federal do Rio Grande do Sul, University of Pennsylvania , Argonne National Laboratory, Ohio State University, Texas A&M University, University of California Santa Cruz, SLAC, Stanford University, Universitäts-Sternwarte München, Ludwig-Maximilians Universität, Excellence Cluster Universe, ETH-Zuerich

- Uncover the physics driving the acceleration of the Universe
- Combine multiple probes of dark energy or modified gravity (galaxy distribution, gravitational lensing, galaxy clusters, supernovae)
- DES is “Stage III”; will be followed by LSST and possibly a spectroscopic survey (both “Stage IV”), extending out to 2030
- 5-year Survey, obtained first light in September, 2012

Dark Energy Survey



- The Dark Energy Camera will take ~400 GByte images per night for 500 nights (200 TeraBytes)
- Will produce catalogs of stars and galaxies along with their properties
- Data Management done at NCSA

Dark Energy Survey Science Program

ANL (Habib), Chicago (Kravtsov), FNAL (Dodelson), SLAC (Wechsler)...

- Analysis tools to be used by the entire collaboration to extract science from the output of Data Management
- Our present focus is getting our codes up and running to be ready for first season of data
- By 2017 we expect to have generated 5-10 Petabytes of simulations; processed all exposures to extract shape measurements; have a working software framework that will combine all sets of observations in a wide variety of ways (i.e. many MCMC's)

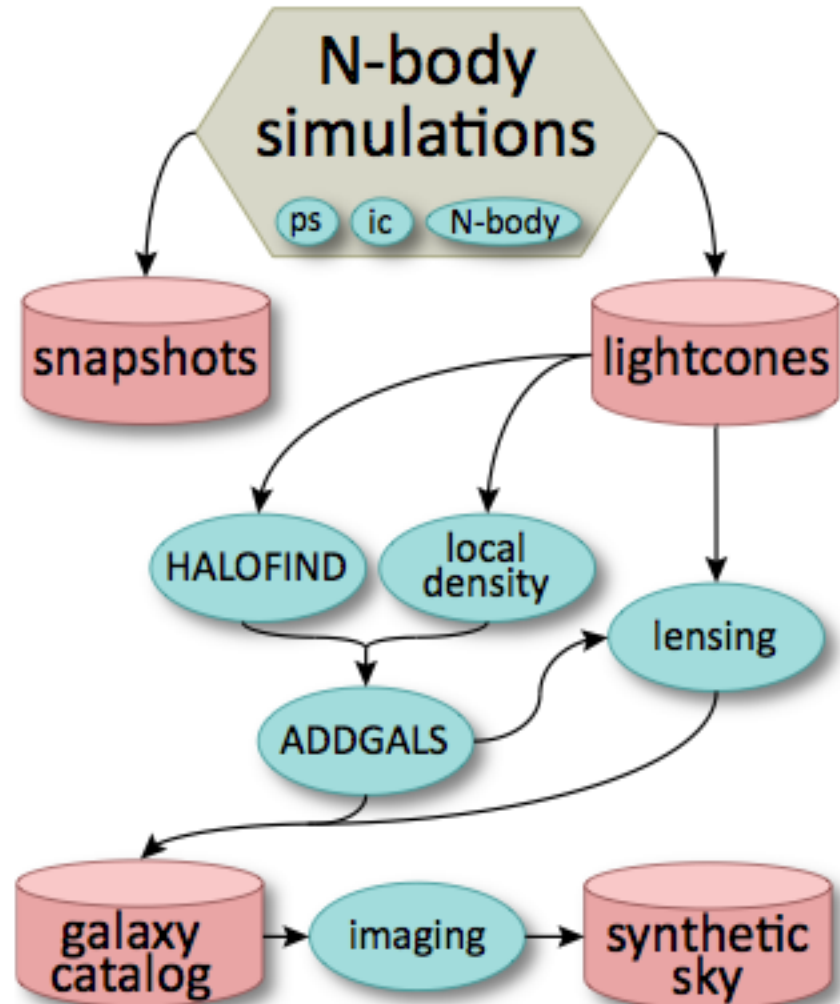
Computational Thrusts

- Simulations
 - N-Body simulations of growth of structure under gravity
 - Halo finding, painting on realistic galaxy distributions, ray tracing
- Value-Added Catalogs
 - Cluster Finders
 - Shape Measurements
 - Galaxy Masks
- Analysis
 - N-point functions
 - Monte Carlo Markov Chain samplers
 - Likelihood codes (heaviest component is theory prediction)

DES-Size Simulations

N-Body simulations are most expensive.

- Currently running several a year, expect to ramp up to 30 by 2017. 0.6M CPU hours/sim
- Use 60k cores with 2-4 GB/core depending on resolution.
- Saving time slices from each N-Body sim → 55TB per sim.
- All of these need to be saved (to re-run post-processing which is constantly evolving). Expect >5TB to be archived by 2017.
- Issue at NERSC might be clock time: Currently runs take 10 days at XSEDE.



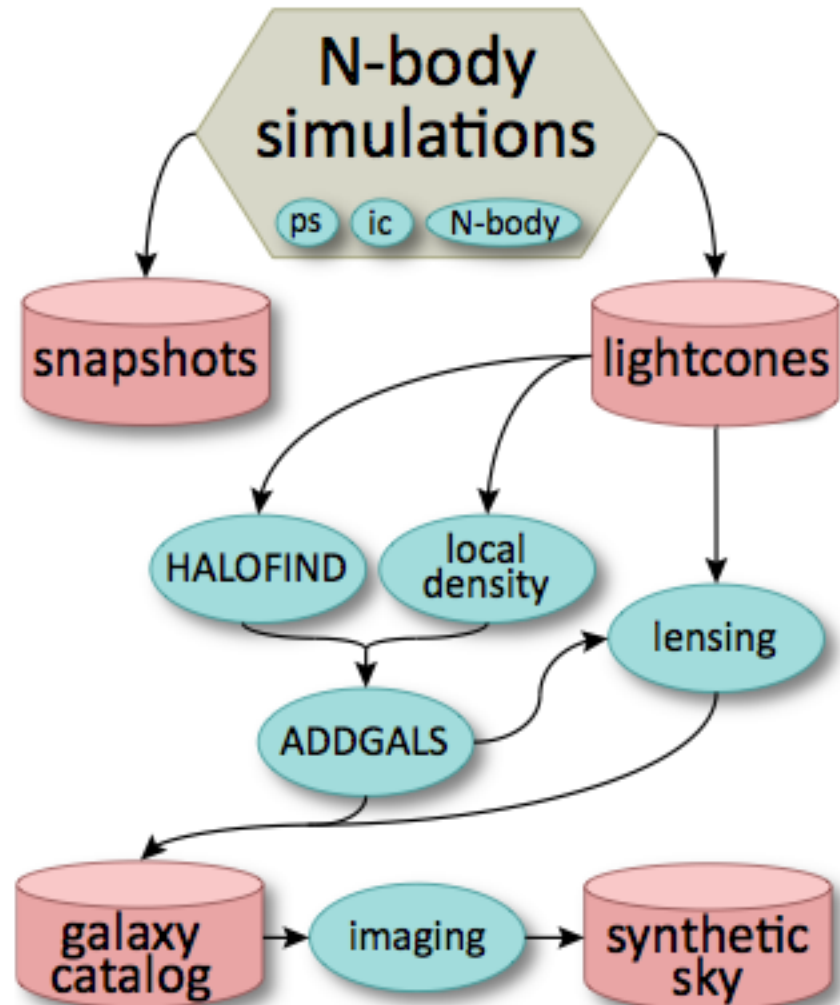
DES Science Program: 2017

	Million CPU Hours	Number of Cores	Memory per Core	Disk Space (TB)	Tape Space (TB)	Number of Runs
Simulations	20	60000	2-4 GB	55	5500	30
Value-Added Catalogs						
Analysis						

DES-Size Simulations

Post-processing done at SLAC.

- Most expensive pieces are **lensing** (ray tracing) and **ADDGALS** (adding galaxies to dark matter halos). Half the CPU time as N-Body runs.
- Currently use 250 cores with 4 GB/core but expect scaling to 2000 cores to be straightforward.
- Halo finder and ADDGALS are embarrassingly parallel. Ray tracing uses MPI.
- Final catalogs ~5 TB → 500 TB by 2017.
- Issue for NERSC: porting code



DES Science Program: 2017

	Million CPU Hours	Number of Cores	Memory per Core	Disk Space (TB)	Tape Space (TB)	Number of Runs
Simulations	20	60000	2-4 GB	55	5500	30
	10	200-2000	4 GB	500	-	100
Value-Added Catalogs						
Analysis						

Value-Added Catalogs

Shape measurements most expensive, currently running at BNL

- Two components: Co-adding single exposures and using point-like stars to estimate instrumental/atmospheric (PSF) effects.
- Plan to run ~25 times/year; each run uses all data taken to date (~130TB) → 32 GB/node
- Embarrassingly parallel, running on 1000 cores
- Each run takes ~0.2M CPU hours → ~10 days
- Issues for NERSC: manpower. Erin Sheldon currently trying to use code at NERSC



DES Science Program: 2017

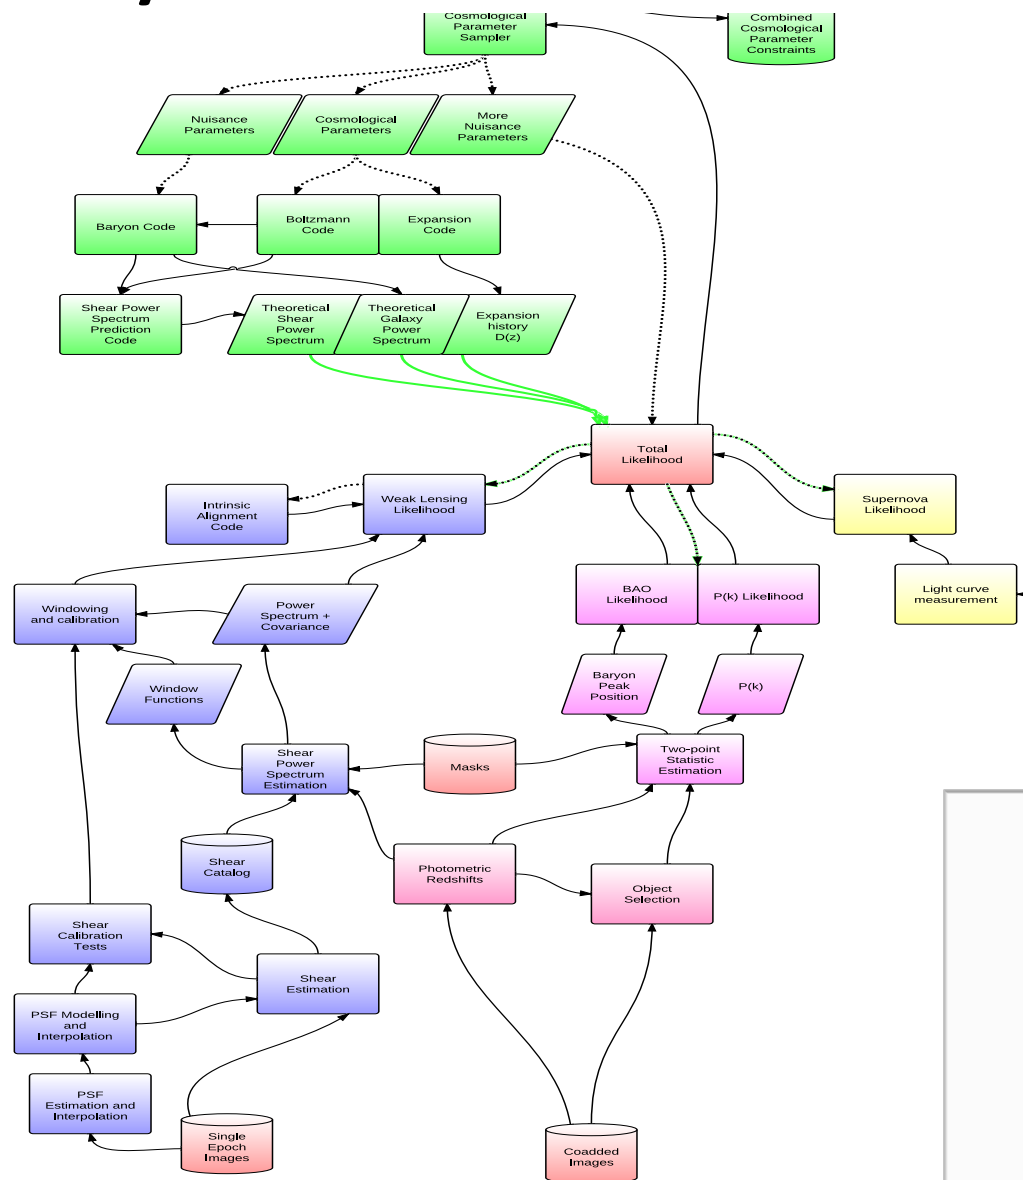
	Million CPU Hours	Number of Cores	Memory per Core	Disk Space (TB)	Tape Space (TB)	Number of Runs
Simulations	20	60000	2-4 GB	55	5500	30
	10	200-2000	4 GB	500	-	100
Value-Added Catalogs	5	1000	32 GB/ node	150	-	25
Analysis						

Analysis

Analysis can be broken down into MCMC sampler ($\sim 5 \times 10^5$ calls) to likelihood code (~ 1 -10 seconds)

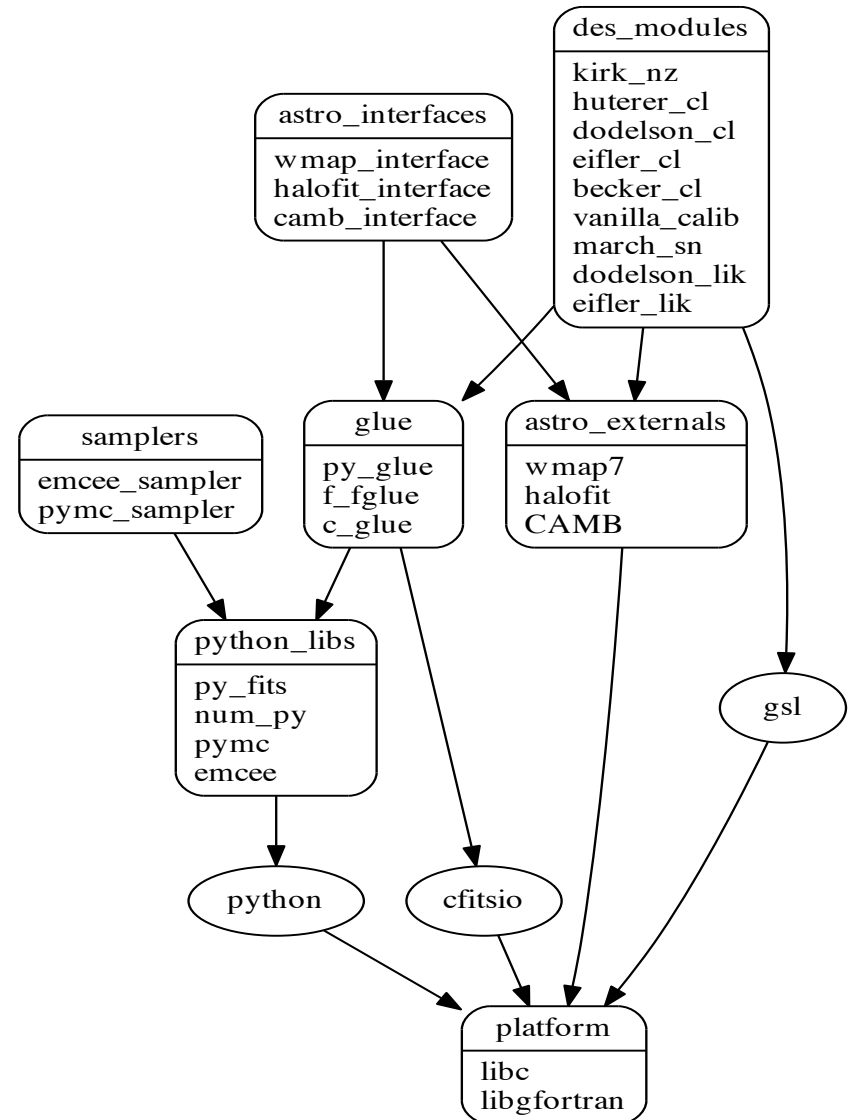
Issues:

- **Combining probes** (expect many different combinations, each with its own covariance matrix) and likelihood
- **Many parameters** (not just cosmology, also nuisance parameters ... and need to unify treatments across probes)
- **Unknown model** (we don't know what is driving acceleration!)
- **Not linear:** cannot simply run *camb*
- **Multiple versions** of each module
- **Many runs:** $\sim 25k$ (100 catalogs x 5 cuts x 10 2-point functions x 5 models)
- **Not power users**



Analysis

Need to worry about dependencies. Will need to regular monitor and install updates to, e.g., camb. If python version changes or camb changes, How does this affect our results? Need tracking mechanism



DES Science Program: 2017

	Million CPU Hours	Number of Cores	Memory per Core	Disk Space (TB)	Tape Space (TB)	Number of Runs
Simulations	20	60000	2-4 GB	55	5500	30
	10	200-2000	4 GB	500	-	100
Value-Added Catalogs	5	1000	32 GB/node	150	-	25
Analysis	10	?	-	-	-	10,000